



ELSEVIER

Cognition 69 (1999) 243–265

---

---

COGNITION

---

---

## Dr. Angry and Mr. Smile: when categorization flexibly modifies the perception of faces in rapid visual presentations

Philippe G. Schyns\*, Aude Oliva\*

*Department of Psychology, University of Glasgow,  
56 Hillhead Street, Glasgow G63 9RE, UK*

Received 27 January 1997; accepted 2 November 1998

---

### Abstract

Are categorization and visual processing independent, with categorization operating late, on an already perceived input, or are they intertwined, with the act of categorization flexibly changing (i.e. cognitively penetrating) the early perception of the stimulus? We examined this issue in three experiments by applying different categorization tasks (gender, expressive or not, which expression and identity) to identical face stimuli. Stimuli were hybrids: they combined a man or a woman with a particular expression at a coarse spatial scale with a face of the opposite gender with a different expression at the fine spatial scale. Results suggested that the categorization task changes the spatial scales preferentially used and perceived for rapid recognition. A perceptual set effect is shown whereby the scale preference of an important categorization (e.g. identity) transfers to resolve other face categorizations (e.g. expressive or not, which expression). Together, the results suggest that categorization can be closely bound to perception. © 1999 Elsevier Science B.V. All rights reserved.

---

### 1. Introduction

Even casual observers would have no difficulty in placing the two face pictures of Fig. 1 in a number of different categories. The top picture might be recognized as a face, as a female, as a young Caucasian face, as a non-expressive face, or as 'Zoe', if this was her identity. In contrast, the bottom picture could be classified as 'John',

\* Correspondence can be sent to either author.

E-mail: philippe@psy.gla.ac.uk

E-mail: aude@psy.gla.ac.uk

0010-0277/99/\$ - see front matter © 1999 Elsevier Science B.V. All rights reserved.

PII: S0010-0277(98)00069-9

who is a male, comparatively older, and apparently angry. These distinct judgments of similar images reveal the impressive versatility and specialization of face categorization mechanisms (e.g. Etcoff and Magee, 1992; Calder et al., 1996). That is, people can make judgments of gender, age, expression, and (if they know the person) identity, based on the same visual input.

Different face judgments, like most other object categorizations, tend to require different information from the visual input. For example, whereas skin texture could reliably indicate age, the diagnostic shape of the mouth (among other relevant features) would be diagnostic of a judgment of expression. Such flexibility, combined with the high efficiency of face categorizations, suggests that human visual processes have developed particularly effective means for the perceptual encodings of faces.

This raises the general issue of the relationship existing between flexible visual categorizations (i.e. the diagnostic use of visual information) and the perception of the stimulus itself. Are they independent, with categorization operating late, on an already perceived input, or are they intertwined, with the act of categorization influencing the early perception of the stimulus?

Most theories of categorization and recognition have neglected this issue even though their mechanisms presume a flexible use of visual information (see Schyns, 1998 for discussions). However, there is evidence and arguments for the view that cognition does not start where perception ends (i.e., that categorization is intertwined with perception, see Schyns et al., 1998 for a review; though see also Pylyshyn, in press). The experiments reported in this article will examine, in the circumscribed domain of faces, the general claim that the selection of information for categorization can modify the perception of the input.

To illustrate the point, look again at the faces of Fig. 1, then blink, squint, or defocus. Other faces should replace those you initially perceived. If this demonstration does not work, step back from Fig. 1 until your perceptions reverse and you see the angry man in the top picture, and the woman with a neutral expression in the bottom picture.

The pictures of Fig. 1 illustrate *hybrid stimuli* (Schyns and Oliva, 1994); a *hybrid face* simultaneously presents two faces, each associated with a different spatial scale. In Fig. 1, fine scale information (more precisely, high spatial frequencies; HSF) represents a non-expressive woman in the top picture and an angry man in the bottom picture. Coarse scale information (i.e. low spatial frequencies; LSF) represents the opposite, i.e. an angry man in the top picture and a non-expressive woman in the bottom picture. A hybrid face therefore dissociates the face information represented in two distinct spatial frequency bandwidths.

This dissociation suggests a method to explore the influence of categorization tasks on scale perception. Suppose you were instructed to judge the expression of each picture of Fig. 1 and that your responses were 'neutral' and 'angry' for the top and bottom faces. From these, we could infer that your categorizations involved fine-scale cues as the use of coarse cues would have evoked 'angry' and 'neutral', respectively. Suppose you were later instructed to identify the same pictures. If your judgments were, from top to bottom, 'John' and 'Zoe' (rather than 'Zoe' and 'John'),

we would infer that coarse scale cues were the bases for your decisions. Such changes in scale preferences would indicate that the information required for these two categorizations can reside at a different spatial scale of an identical picture, but also that your perceptual systems flexibly tuned to preferentially extract information at these scales.

At an empirical level, the experiments reported here used hybrid stimuli to demonstrate that different face categorizations can flexibly use the information associated with different spatial scales, even when relevant cues are available at the other scale. Face stimuli offer advantages over other objects and scenes for this demonstration: Their compactness enables a tight control of presentation which limits the scope of useful cues; the familiarity of their categorizations (e.g. gender, expression, identity) simplifies the experimental procedure which does not require prior learning of the multiple categories: most people are ‘natural’ face experts (Bruce, 1994).

At a more theoretical level, we sought evidence that a diagnostic use of spatial scales in a categorization task can modify the immediate perceptual appearance of the input (Schyns, 1997). Spatial scales have a privileged status in perceptual organization: they are known to support a wide range of low-level visual tasks, including motion (Morgan, 1992), stereopsis (Legge and Gu, 1989), edge detection (Marr and Hildreth, 1980), depth perception (Marshall et al., 1996) and saccade programming (Findlay et al., 1993). Thus, evidence that different face categorizations modify the perception of spatial scales would suggest that the diagnostic use of visual information interacts with the early stages of visual processing.

In sum, we are studying the theoretical issue of whether constraints from categorization can change the use and perception of spatial scales. We are not addressing the problem of face perception *per se*, but are using face stimuli because they are more convenient than other objects or scenes.

The following section briefly reviews the notions of simultaneously filtering the input at multiple spatial scales, of category specific cues at different scales, and it examines how the diagnostic use of these cues could interact with scale perception.

### *1.1. Scale perception and scale-based recognition*

It can be shown from Fourier’s theorem that a two-dimensional signal can be decomposed into two sums of sinusoids (with different amplitudes and phase angles), which each represents the image at a different spatial scale. Psychophysical studies on contrast detection and frequency-specific adaptation revealed that our perceptual system analyzes the visual input at multiple scales (see De Valois and De Valois, 1990, for a review) through banks of independent, quasi-linear, band-pass filters, each of which is narrowly tuned to a specific frequency band (Campbell and Robson, 1968; Pantle and Sekuler, 1968; Blackemore and Campbell, 1969; however, see also Henning et al., 1975, for signs of interactivity and Snowden and Hammett, 1992, for non-linearity). It is now generally acknowledged that spatial filtering is a common basis for the extraction of visual information from luminance contrasts (gray-levels; Marr and Hildreth, 1980).

It is worth pointing out that the temporal integration of spatial scales is very fast, and that their detection does not vary appreciably over a wide range of spatial frequencies at very brief presentations (e.g. 30 ms, Hammett and Snowden, 1995). However, different spatial scales tend to convey different information for recognition (e.g. Bachmann, 1991; Parker et al., 1992, 1996; Costen et al., 1994, 1996; Schyns and Oliva, 1994; Oliva and Schyns, 1995, 1997; Hughes et al., 1996; Dailey and Cottrell, in press). For faces, HSF tend to represent fine scale cues such as age and expression-related wrinkles, eyelashes, the precise contours of the nose, the eye, the mouth, the chin and other sharp face boundaries. LSF represent coarse blobs which, individually, do not have sufficient resolution to be identified as a nose or an eye, but which, together, represent the informative configural structures of faces (as revealed in Fig. 1). Thus, the discovery that a limited bandwidth of spatial frequencies is sufficient to resolve a face categorization indicates that this bandwidth mediates some of the visual cues that could support this task.

Work in scale-based recognition of coarse-quantized face pictures<sup>1</sup> has revealed that a critical bandwidth of LSF (between 8 to 16 cycles/face) can mediate their identification (i.e. with a performance superior to 80%, see Bachmann, 1991; Costen et al., 1994, 1996). LSF information can therefore mediate face identification<sup>2</sup>. Expertise with faces tends to induce the development of configural and global, rather than componential and local, recognition strategies (e.g. Carey, 1992; Tanaka and Farah, 1993; Tanaka and Sengco, 1997; see Diamond and Carey, 1990; Tanaka and Gauthier, 1997, for discussions of extensions of this expertise to other stimuli). Configural properties tend to be better represented at coarser scales (Bachmann, 1991; Costen et al., 1996), and a LSF (instead of a HSF) bias might therefore be expected from face experts resolving speeded identification tasks. Dailey and Cottrell (in press) reported such LSF preference in a neural network (a *mixture of experts*, Jacobs et al., 1991) in which two modules competed to identify 12 faces. One module was only provided with LSF information whereas the other module received the complementary HSF. In agreement with the development of expert configural identification strategies, the LSF-fed module won the competition.

There is comparatively less research on the spatial scales supporting other face categorizations than identity. Sergent (1986) and Sergent et al. (1992) suggested that intermediate spatial scales could underlie gender judgments, a simple task that a

<sup>1</sup>The famous quantized portrait of Abraham Lincoln illustrates a quantized image (see Harmon, 1973). The original image is typically low-passed filtered. It is then divided into a number of blocks of equal size. Within each block, pixel intensities are averaged. These images are then described in terms of the maximum number of block alternations in the image (cycles/image). For example, a  $256 \times 256$  face image quantized at 8 cycles/face comprises 16 blocks of 16 pixels per axis.

<sup>2</sup>Precisely which LSF, however, is difficult to assess because the quantization procedure has a number of shortcomings: it introduces HSF masking noise (corresponding to the mosaic of fine scale edges representing the blocks) and it also distorts the configurations represented in LSF (see Bachmann, 1997, for a discussion of the shortcomings of quantized faces). We did not use quantized faces in our experiments.

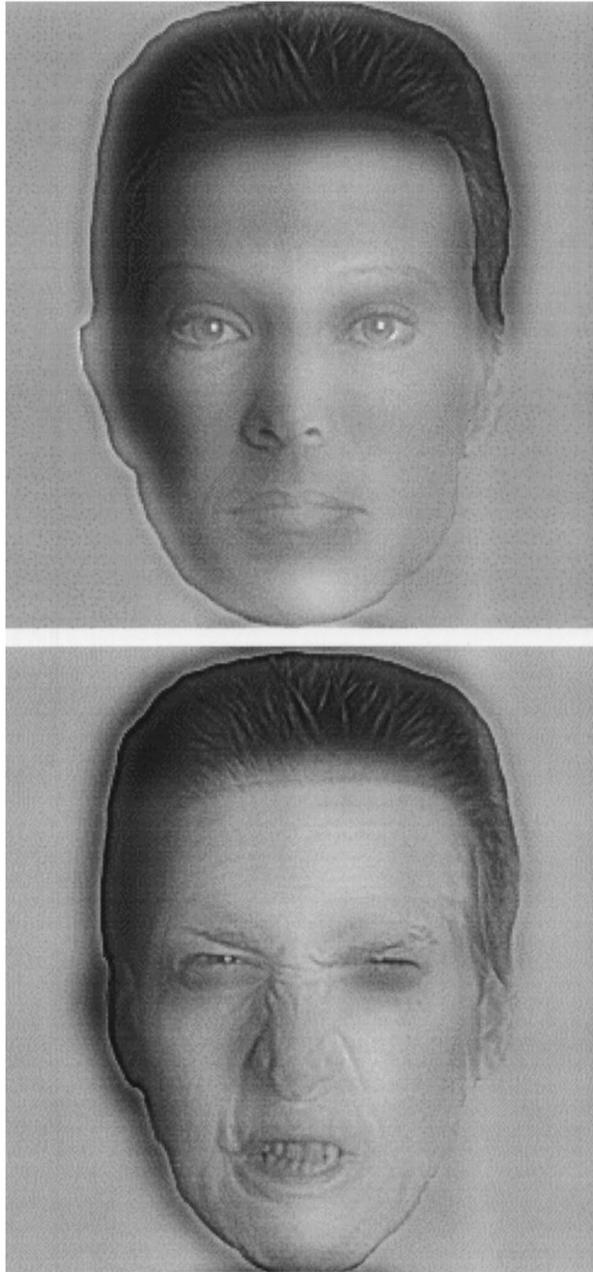


Fig. 1. Two of the hybrid faces used in experiments 1, 2 and 3. The fine spatial scale (HSF) represents a non-expressive woman in the top picture and an angry man in the bottom picture. The coarse spatial scale (LSF) represents the angry man in the top picture and the neutral woman in the bottom picture. To see the LSF faces, squint, blink, or step back from the picture until your perception changes.

perceptron can already solve (Gray et al., 1995). Expressions, however, have recently come under closer scrutiny in the literature (Calder et al., 1996). The early study of Eckman and Friesen (1975) explored universal cues that could underlie the categorization of six facial expressions (happy, angry, sad, surprise, fear and disgust). These cues were all local and their best support were HSF (see Rosenblum et al., 1996 for a model of facial expression based on HSF). However, Bassili (1979) suggested that two blobs around the mouth (a coarse scale cue) could subtend the detection of the happy expression, and developmental studies also suggested that coarse information was sufficient for the detection of expressions in 3-month-old babies (Barrera and Maurer, 1983; Kuchuk et al., 1986).

In summary, there is suggestive evidence that information at different scales can support different categorizations of one face. However, the question remains of whether information at any of these scales (e.g. LSF for the identity of faces) is *selectively* accessed and used when relevant cues are also present at the other scale (e.g. HSF), as they are with all naturalistic stimuli. In the face studies reviewed so far, the evidence that information at a restricted bandwidth is *sufficient* for this or that categorization does not imply that it is selectively accessed when all spatial scales are available. Studies with hybrids, however, circumvent this difficulty because (1) hybrids are full-bandwidth stimuli which do not limit the available scale information, and (2) their LSF and HSF represent opposite face information (e.g. LSF male vs. HSF female, or LSF John vs. HSF Lynda). Thus, evidence of categorization at a single bandwidth not only indicates that its information is sufficient for the task at hand, but also that this scale is preferred over other representations of the relevant information at other, competitive scales. Band-filtered stimuli or quantized pictures do not warrant this conclusion because relevant information is present at only one scale and so information from different scales does not compete as in hybrids.

Still, any study of selective use would be severely limited if, as is often assumed, early vision was organized to always process coarse scales before fine scales (e.g. Breitmeyer and Ganz, 1976; Parker et al., 1992; Schyns and Oliva, 1994; Hughes et al., 1996). Oliva and Schyns (1997) argued against such mandatory perceptual determination in favor of a more flexible processing. In their Experiment 3, two groups were initially trained to categorize hybrid scenes which comprised relevant cues at only one spatial scale; the other scale was noise (either LSF or HSF, depending on group). Subsequent to this sensitization, and without any discontinuity in the presentation of stimuli, subjects categorized hybrids that comprised a different scene at each of the two scales. Results showed that subjects maintained their categorizations at (and were only aware of) the scale congruent with their sensitization stimuli (LSF or HSF). This experiment suggested that rather than a mandatory use of one scale before the other, people can selectively attend to and recognize the information of different spatial scales. In Oliva and Schyns (1997), however, flexible scale perceptions resulted from initial exposure to LSF vs. HSF stimuli, not from a different categorization task (e.g. identity vs. facial expressions) on the same stimuli. The latter is more general because it isolates a top-down, not a bottom-up, determination of flexible use and perception of scales.

This paper reports three experiments which examine the general claim that categorization tasks differ in their scale use and perception. Experiment 1 required subjects to categorize face hybrids similar to those of Fig. 1 according to whether they were male vs. female, expressive vs. non-expressive, and their specific face expression. We observed how the information required by these common categorizations affected scale biases and perception. Experiment 2 examined whether the scale bias needed to resolve a first task transferred across tasks to bias a second categorization. Experiment 3 applied these principles to understand the bias that arises when people initially learn to identify faces, arguably their most important categorization, before resolving other tasks (i.e. gender, expressive, and which expression). In each experiment, the emphasis was on providing an ‘existence proof’ that categorization can determine a flexible use and perception of spatial scales, using faces as a case model.

## 2. Experiment 1

Experiment 1 tests the hypothesis that vision can selectively use the different spatial scales of an identical face hybrid based on the information requirements of different categorizations. Each hybrid combined a man and a woman, only one of which was expressive (happy or angry), the other one was neutral (as in Fig. 1). Consequently, three distinct categorizations could be applied to the same stimulus: male vs. female, expressive vs. non-expressive, and happy vs. angry vs. neutral. Because each scale in a hybrid represents a different categorization response, the latter could be used to assess the scale biases of the categorizations.

As explained earlier, hybrids are well-suited to study scale biases because they create a cue-conflict situation but are themselves unbiased (relevant information resides at both scales). To tap into the early selection of visual information, it is necessary to adjust to its own time course. Psychophysical research has demonstrated that spatial scales are integrated very early, within the first 50 ms from stimulus onset. In recognition, Oliva and Schyns (1997, Experiment 1) showed that a 30 ms, masked presentation of a single hybrid scene successfully primes the recognition of the two scenes (LSF and HSF) it represents. Hence, our studies will require the very brief, tachistoscopic presentations of stimuli that are ubiquitous in psychophysics. In fact, these conditions have already elicited selective perception and use of scales (Oliva and Schyns, 1997). Our methods are therefore adjusted to the theoretical issue of whether the act of categorization can itself induce selective use and perception of spatial scales.

### 2.1. Method

#### 2.1.1. Subjects

Forty-five University of Glasgow students (male and female between 17 and 27 years of age) with normal or corrected vision were paid to participate in the experi-

ment. They were randomly assigned to experimental groups with the constraint that the number of subjects be equal in each group.

### 2.1.2. Stimuli

Stimuli were composed from a face set of six males and six females (all professional actors), each of whom displayed three different expressions (for a total of 36 original face stimuli). Expressions were Happy (H), Angry (A) and Neutral (N). Henceforth, we will call *expressive* the H and A face pictures, and *non-expressive* the N pictures. The faces were professionally photographed in controlled conditions of illumination. Hairstyle is a reliable cue to gender and so it was digitally normalized to the unisex hairstyle that characterizes both pictures of Fig. 1. We also normalized stimulus size to insure that the two face pictures composing each hybrid would overlap. Again, it is worth pointing out that if this changes the natural distribution of face cues, our main goal is not real-world face perception but flexible scale use. We synthesized 112 different hybrids (16 for practice and 96 for testing) by combining the face of a man with the face of a woman, only one of which was expressive, the other one was neutral (see Schyns and Oliva, 1994). Assignment of gender and expression was counterbalanced across spatial frequencies (LSF, below 2 cycles/deg of visual angle; HSF, above 6 cycles/deg of visual angle). These low- and high-pass cut-offs corresponded to SF below and above 8 and 24 cycles/image, respectively<sup>3</sup>. Hybrids were  $256 \times 256$  pixels on a 256 gray-level scale.

The set of 112 hybrids could either be decomposed into 56 LSF-male/HSF-female and 56 LSF-female/HSF-male, or into 56 LSF-expressive/HSF-neutral and 56 LSF-neutral/HSF-expressive. Expressive faces were equally divided between 28 angry and 28 happy in LSF and HSF.

### 2.1.3. Procedure

Each subject was assigned to one of three different categorization tasks. In the GENDER task, they were instructed to decide whether the presented stimulus was male or female. In EXpressive vs. Non-EXpressive (EXNEX), they were instructed to determine whether or not the stimulus was expressive. In CATegorization of EXpressions (CATEX), they were instructed to identify the expression of the stimulus (possible expressions were H, N, A). Thus, each group performed a different task on the same set of 112 stimuli. Subjects were run on small blocks of 56 stimuli at a time. In each block, stimuli were randomized and the first eight served as practice trials. Order of blocks was counterbalanced across subjects.

Stimuli were presented for 50 ms on the computer monitor of an Apple Macintosh PowerPC. Subjects indicated their response by pressing the appropriate keyboard key in the GENDER and EXNEX tasks. In CATEX, they wrote down the first letter corresponding to the expression they perceived (either N, H or A). Instructions emphasized speeded decisions, but they did not disclose the ambiguity of the

<sup>3</sup>The 'blocked face' literature (Bachmann, 1991; Costen et al., 1994, 1996) suggests that there is enough LSF information at 8 cycles/image to obtain more than 80% identification accuracy, and that 24 cycles/image represent all the boundary edges defining important face components.

hybrids. After the experiment, we assessed whether subjects noticed the ambiguity. The experimenter put a hybrid on the screen then asked the following question: ‘Here is a stimulus composed of two faces (the experimenter would point at the two faces). Did you explicitly notice, or did you have the impression that there were such ambiguous stimuli in the experiment?’ Subjects’ answers were recorded.

## 2.2. Control experiment

As explained earlier, we used subjects’ categorization responses to measure their LSF and HSF scale biases. However, this method is valid only to the extent that each task can be independently solved using the LSF and HSF of the original faces. We run an independent control with 36 subjects divided into three groups. To recreate conditions of stimulation comparable to those of Experiment 1 (i.e. full-bandwidth not band-filtered, nor quantized stimuli), we synthesized hybrids which added noisy features to the scale (either LSF or HSF) that did not represent a face (as in Oliva and Schyns, 1997). These features were randomly chosen face parts (e.g. nose, eye, mouth) placed at randomly chosen locations within the image plane to decorrelate the parts with their expected location (see Fig. 2)<sup>4</sup>.

Each subject group was asked to solve a different task (GENDER, EXNEX, and CATEX) on LSF and HSF noisy hybrids. At 50 ms presentations, gender judgments were 89% correct in LSF and 76% correct in HSF. These percentages were respectively of 81% and 77% for EXNEX, and 81% and 76% for CATEX. Note that although there is a LSF bias in all tasks, they can all be solved with similar error rates on the basis of all spatial scales, validating the method of using subjects’ categorization responses to infer scale usage<sup>5</sup>.

## 2.3. Results and discussion

The data of nine subjects (from a total of 45) were discarded from analysis because they claimed in their debriefing to have perceived two faces in the hybrid stimuli. These perceptions could induce a bias because subjects could strategically decide to only report the face associated with one of the two scales. For the remaining subjects, we subtracted their LSF from HSF categorization percentages and compared this score with 0 (indicating an equal use of LSF and HSF). A significant HSF bias was observed in the EXNEX task (LSF = 38%, HSF = 62%,  $t(11) = 2.38$ ,

<sup>4</sup>We chose not to present LSF and HSF information by itself because band-pass stimuli are not comparable to full-bandwidth stimuli (as hybrids are). See Hughes et al. (1996) for difficulties with using band-passed stimuli.

<sup>5</sup>We designed our experiments to look for qualitative differences in the perception of identical stimuli. We therefore have no hypothesis about the time course of these perceptions, and even if we did, reaction times would not offer the best approach. Categorization responses are typically produced within the first 400 to 1000 ms following stimulus onset. In contrast, scale analysis is known to occur within the first 50 ms of processing. Hence, the time resolution of a categorization task is simply too coarse to reveal something about cognitive influences on the microgenesis of orthogonal scale perceptions.

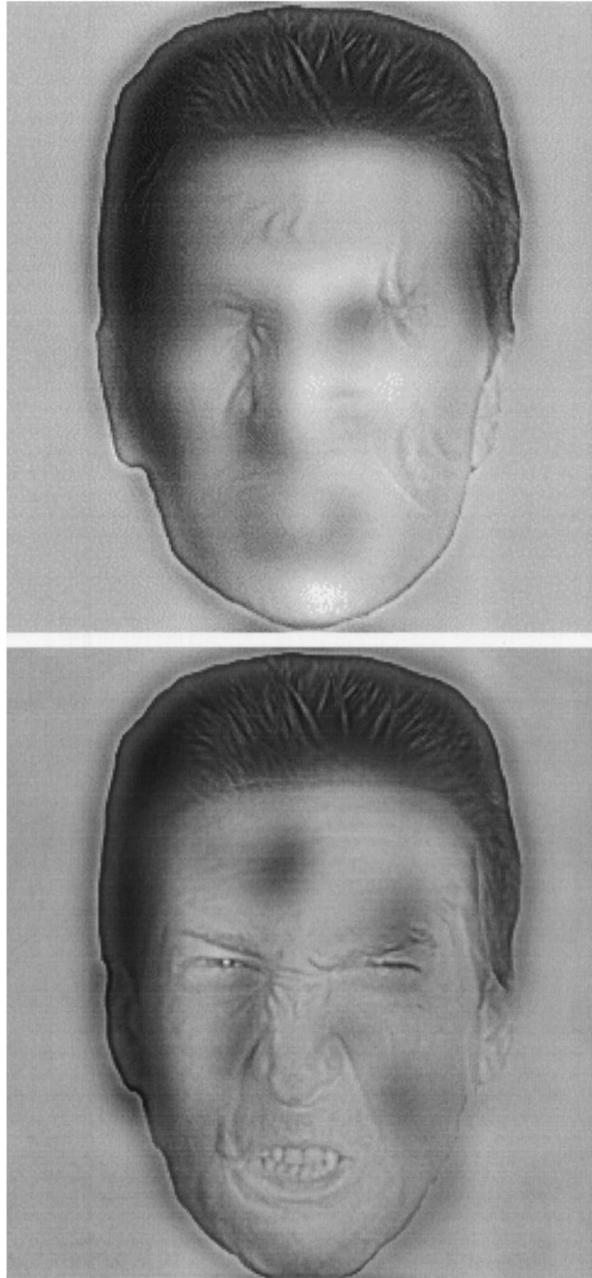


Fig. 2. The hybrid stimuli used in the control of Experiment 1. The top picture represents an angry male in LSF with HSF noise. The noise represents face features at random locations within the face. The bottom picture represents the opposite: the HSF angry male with LSF noise.

$P < 0.05$ ), a significant LSF bias in CATEX (LSF = 66%, HSF = 30%, 4% error<sup>6</sup>,  $t(11) = 3.86$ ,  $P < 0.01$ ), and no bias in GENDER (LSF = 52% and HSF = 48%,  $t(11) < 1$ , see Fig. 3).

In the two biased tasks (EXNEX and CATEX), subjects were individually labeled either as LSF-biased or HSF-biased using a 55% categorization response threshold. In total, eight EXNEX subjects were HSF-biased, with only one LSF-biased and three undetermined. In contrast, nine CATEX subjects were LSF-biased with one HSF-biased and two undetermined. A chi-square test of association revealed that there was a significant association between categorization tasks (EXNEX vs. CATEX) and scale biases (LSF vs. HSF),  $\chi^2(1) = 8.87$ ,  $P < 0.01$ .

Notice that the reported orthogonal scale biases arose from an identical stimulus set. Hence, the argument does not hold that a stimulus difference induced the orthogonal biases. These, in conjunction with the debriefings, suggest that subjects did not notice the scale they did not categorize. In other words, the happy face of one group was perceived as non-expressive in the other group. We can therefore propose that the task changed the perceptual content of an identical stimulus.

The question remains of why CATEX and EXNEX induced different biases. It would seem that a precise judgment of expression should require precise HSF cues, whereas the cruder EXNEX judgment would not need such precision. A careful examination of the task<sup>7</sup>, however, reveals that EXNEX can be resolved by deciding whether the lips are straight (a HSF local information), whereas if a decision of expression required an analysis of the entire face, LSF would be the better information with the fast presentations used here. We will not dwell too much on this interpretation here, as our main purpose is not to study the specific cues supporting different face categorizations (though see Jenkins et al., 1997).

### 3. Experiment 2

Experiment 1 showed that categorizations can differently bias the scale processing of stimuli. Interestingly, the three tasks of Experiment 1 induced the three possible instances of the bias: No bias in gender, a HSF bias in expressiveness and a LSF bias in categorization.

One question that arises is the perceptual status and the dynamics of scale biases driven by the search for diagnostic information: Are they strictly transient and local to the task at hand, or do they possess a ‘perceptual inertia’ that transfers across tasks? In analogy to the ‘problem-solving set’ in which one encoding of a situation blocks another successful encoding (Mayer, 1992), a ‘perceptual set’ rigidity could be demonstrated in which the scale bias of a first categorization subsequently biases

<sup>6</sup>Error means choosing an expression which was neither represented in LSF nor in HSF.

<sup>7</sup>The notion of task here encompasses not only the categorization instruction, but also the actual stimulus set, its associated distribution of categorization-relevant cues, and the specifics of stimulus presentation.

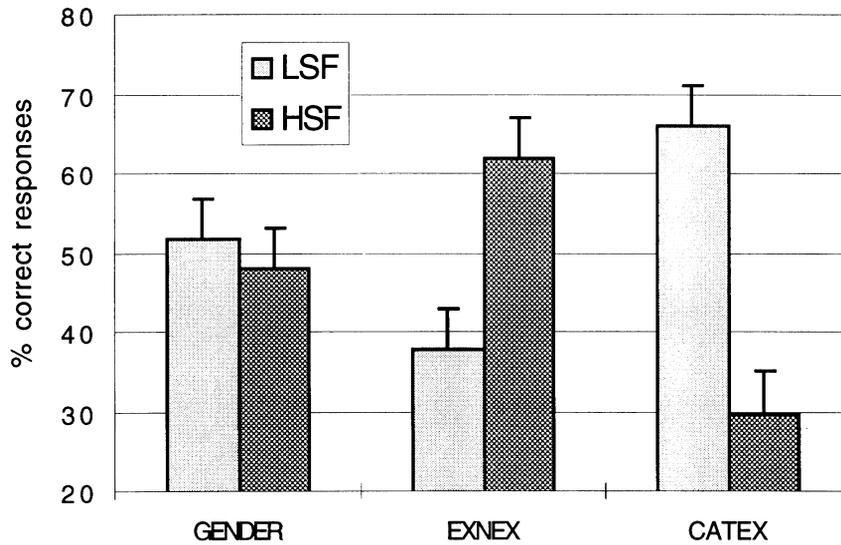


Fig. 3. The scale biases of the three different categorizations (EXNEX, CATEX and GENDER) in Experiment 1. The Figure illustrates that different biases were obtained for different categorizations of the same hybrid pictures.

the perceptual encoding of a second categorization (see, e.g. Long et al., 1992, for evidence of set biases on the perception of the Necker cube). Lasting perceptual transfers would confirm the perceptual nature of the categorization-induced biases. Transfer would be particularly important if one prominent categorization (e.g. identification in the case of faces) modified the scale at which other categorizations are performed (e.g. expression, gender). Experiment 2 sets the stage for Experiment 3 where this issue is explicitly addressed.

In Experiment 2, two groups were initially asked to resolve a different categorization of the same hybrid faces. One group was assigned to the EXNEX task (which from Experiment 1 is known to induce a HSF bias) whereas the other group solved the CATEX task (which induces a LSF bias). Following this, subjects were asked to determine the gender of hybrids (an unbiased task). We then observed whether the orthogonal biases acquired in the initial categorization would transfer to the gender task (unbiased in Experiment 1) to induce orthogonal perceptions of the same stimuli.

### 3.1. Method

#### 3.1.1. Subjects

Thirty University of Glasgow students (male and female between 17 and 27 years of age) with normal or corrected vision were paid to participate in the experiment. They were randomly assigned to an experimental group with the constraint that the number of subjects be equal in each group.

### 3.1.2. Stimuli

Hybrid faces were strictly identical to those of Experiment 1: 112 hybrids divided into 96 test and 16 practice stimuli. Because Experiment 2 comprised two successive tasks, the 112 stimuli were used twice, once per task.

### 3.1.3. Procedure

As in Experiment 1, stimuli were presented one at a time, for 50 ms, on the computer monitor. In the EXNEX group, subjects were instructed to determine whether or not the presented stimulus was expressive. In the CATEX group, subjects were instructed to determine the expression of the stimulus. Following this, subjects were shown the stimuli again, one at a time for 50 ms on the computer screen, but this time they were all instructed to determine their gender. In each task, stimuli were presented in blocks of 48 trials, and the order of blocks was counterbalanced across subjects. At the end of the experiment, we ran a debriefing identical to Experiment 1, to determine whether subjects had perceived the presence of two faces in the experimental materials.

## 3.2. Results and discussion

The data of six subjects (from a total of 30) were discarded from analysis because their debriefing revealed that they perceived two faces in the hybrids. As in Experiment 1, we traced the LSF vs. HSF categorization biases through categorization responses. EXNEX was biased to HSF (HSF = 60%, LSF = 40%, a *t*-test on the difference score between LSF and HSF revealed that this difference was significant,  $t(11) = 4.17$ ,  $P < 0.01$ ) whereas CATEX was biased to LSF (LSF = 64%, HSF = 30%, 6% error,  $t(11) = 6.86$ ,  $P < 0.0001$ ). These biases replicated those observed in Experiment 1.

We can now turn to the main issue of Experiment 2: does the bias acquired in the context of solving a first task transfers to the resolution of a subsequent unbiased task? Analysis of the scale biases in the subsequent GENDER task revealed that this was indeed the case. EXNEX subjects manifested a significant 58% HSF bias for gender, one-tailed  $t(11) = 3.02$ ,  $P < 0.01$ , whereas CATEX subjects were 79% LSF-biased in the same task, one-tailed  $t(11) = 4.01$ ,  $P < 0.01$ . The scale biases were therefore mutually exclusive.

It is important to emphasize that in Experiment 2, both the stimuli *and* the categorization task (GENDER) were identical across groups in the testing phase. This means that the same stimulus was more likely to be perceived as a female in one group, and as a male in the other group, depending on previously acquired biases. In other words, the perceptual set effect induced the perception of different contents in identical stimuli. This set effect is different from the typical sensitization to different stimuli reported in the literature. For example, in Long et al. (1992), one participant group was initially exposed to pictures representing one interpretation of the bi-stable Necker cube while the other group saw pictures representing the other, mutually exclusive interpretation. The groups were then transferred to the famous ambiguous figure and their perceptions of this stimulus were mutually exclusive. Our results

share these mutually exclusive perceptions of hybrid faces, but our sensitization stimuli were identical in the groups—i.e. we did not initially present one group with LSF representations of faces, and the other group the complementary HSF representations, but see Oliva and Schyns (1997) for this procedure applied to scenes.

In sum, the constraint of locating relevant information in a first task set perceptual systems to differently encode and perceive identical stimuli in the following task. This transfer illustrates one form of perceptual changes that can accompany multiple categorizations of stimuli.

#### 4. Experiment 3

The transfer of scale bias from an initial categorization to a subsequent task raises new issues. For example, identification is a particularly important face categorization which could bias the perceptual encodings of other tasks, possibly overriding their ‘spontaneous’ bias, i.e. the bias that occurs when the identity of faces is not known (as in experiments 1 and 2). The cues of identity would then also support gender, expressiveness and expression.

In Experiment 3, subjects first learned to expertly identify (i.e., without a single mistake) the pictures of six faces (three males and three females undergoing three different expressions). These faces composed hybrids in a second phase to assess the scale preference of identification. In a third phase, subjects were divided into six groups. Three of them solved the GENDER, EXNEX and CATEX tasks using hybrids made from the learned faces. The remaining groups solved the same tasks using hybrids made from unknown faces. We first observed whether expertise with faces induced a perceptual set which biased subsequent categorizations. A further comparison between the conditions of known and unknown test faces isolated possible effects of identity on the scale biases of GENDER, EXNEX and CATEX.

##### 4.1. Method

###### 4.1.1. Subjects

Seventy-two University of Glasgow students (male and female between 17 and 27 years of age) with normal or corrected vision were paid to participate in the experiment. They were randomly assigned to condition with the constraint that the number of subjects be equal in each group.

###### 4.1.2. Stimuli

Stimuli were hybrid faces. The 12 original faces were split into two subsets with equal numbers of males and females. As there were three original pictures per face (one per expression), each subset comprised a total of 18 pictures from which we computed 72 hybrids as in Experiment 1. We systematically combined different identities in LSF and HSF, only one of which was expressive: 6 faces  $\times$  3 faces of the other gender  $\times$  2 expressions  $\times$  2 spatial scales = 72. The two non-intersecting subsets of 72 hybrids were used as different testing stimuli.

#### 4.1.3. Procedure

To gain expertise with faces, subjects went through an extensive learning procedure distributed over several days.

*4.1.3.1. Initial learning phase.* In the initial phase, subjects were instructed to learn the name associated with the original face pictures. Half of the subjects learned the identities of the first face subset (which were arbitrarily called *Mike, Peter, Simon, Helen, Linda* and *Mary*); the second half learned the identities of the other subset (called *Edward, John, Tony, Anne, Jenny* and *Patricia*). Each face picture was glued on cardboard with its name appearing at the bottom.

Subjects were told to disregard the haircut (which was normalized across pictures) and to learn the name-picture pairings at their own pace, until they felt they knew them. This took typically between 15 to 30 min. Subjects then saw the 18 face pictures, one at a time, without their names, and their task was to identify them aloud. A single mistake would restart the self-paced learning procedure until subjects could correctly identify all pictures.

*4.1.3.2. Second learning phase.* Subjects were asked to come back a couple of days later to crystallize their learning. They first repeated one learning stage as just described. One hundred percent identification accuracy was again required before subjects could move on to the next learning phase. This consisted of randomly presenting the 18 original face pictures, one at a time, on the computer screen, for 50 ms. Subjects had to write down the name associated with the face and then check whether this was correct by pressing the computer keyboard space-bar. Following four such blocks of 18 training stimuli, identification performance was measured on two successive blocks. A 100% identification accuracy was required before subjects moved on to the next stage. Five (of 72) subjects could never reach this level of performance and stopped the experiment at the second learning phase.

*4.1.3.3. Identification test.* Subjects who reached this test were all expert identifiers of 6 faces. Their scale biases were assessed using the 72 hybrids computed from the original 18 face pictures they knew. Hybrids were presented one at a time, on a computer screen, for 50 ms. Subjects were told to identify the presented picture (they were not told these were hybrids) and write down its name on the provided answer sheet (in a six-alternative-forced-choice paradigm—an error here is naming a face not represented in either the LSF or the HSF of the hybrid). Of the remaining 67 subjects, eight did not reach a minimum of 70% identification accuracy and stopped the experiment at this stage. The intention was that subjects not only had sufficient expertise with their faces, but that they could also adequately identify them in hybrids.

*4.1.3.4. Transfer test.* This phase appraised the perceptual set of identity when solving the GENDER, EXNEX and CATEX categorizations. Subjects were here split into six groups, each one of which only resolved one task. Three groups (U\_GENDER, U\_EXNEX and U\_CATEX) resolved either GENDER, EXNEX or

CATEX on hybrids composed of Unknown faces, to isolate the scale bias of the initial identification task. The other three groups (K\_GENDER, K\_EXNEX and K\_CATEX) performed the same categorizations (GENDER, EXNEX or CATEX) on faces they already knew. The tested faces were the same across the six groups; only their status of unknown vs. known differed across the groups (U\_GENDER, U\_EXNEX and U\_CATEX vs. K\_GENDER, K\_EXNEX and K\_CATEX). This enabled a measure of any supplementary effect of knowing the faces while keeping the stimulus base constant across groups.

*4.1.3.5. Debriefing phase.* Following the transfer phase, subjects were shown a hybrid picture and were asked whether they noticed that the experimental stimuli comprised two faces. Their answers were recorded.

## 4.2. Results and discussion

A total of 11 subjects (from the remaining 59) noticed that two faces composed the hybrids; their data were discarded from the analysis. All remaining 48 subjects (eight per group) could accurately identify the faces they learned without ever noticing that two of them systematically composed the hybrids.

Hybrids were correctly identified (i.e. either on the basis of LSF or HSF) on 82% of the trials. Subjects were able to correctly identify the stimuli on the basis of their LSF and HSF, but the former dominated at 90%, with no difference between the two face sets (LSF = 91% vs. 89%). This is in line with the literature on quantized faces reported earlier which showed that a eight cycles/image threshold (which is equal to our 2 cycles/deg LSF filtering) can support identification (Bachmann, 1991; Costen et al., 1994, 1996). Evidence of LSF supporting identity must always be interpreted in the context of the number of faces to identify. If LSF simply represented degraded information then it would be likely that supplementary HSF would be needed to identify a larger number of faces. However, if LSF also represented a unique configuration of the faces, then it could suffice for identification. A parametric study relating numbers of faces to sufficient scale information for their identification is outside the scope of this paper, but it should be carried out.

The LSF identification bias transferred to all other tasks (see Fig. 4). In comparison to experiments 1 and 2, it is interesting to note a reversal of bias from HSF to LSF in EXNEX. Note that familiarity with the stimuli did not cause the reversal because it is also observed in the groups which did not know the tested faces (the U\_\* groups).

A two-way, between-subject ANOVA (known vs. unknown test faces  $\times$  GENDER, EXNEX, CATEX) on the difference scores between LSF and HSF categorizations revealed significant main effects of face knowledge,  $F(1,42) = 8.93$ ,  $P < 0.01$ , categorization task,  $F(2,42) = 18.17$ ,  $P < 0.0001$ , and a significant interaction between the two factors,  $F(2,42) = 3.25$ ,  $P < 0.05$ . Post-hoc comparisons (Newman-Keuls) between known and unknown test faces revealed significant differences in GENDER,  $P < 0.05$ , and CATEX tasks,  $P < 0.05$ , but not in EXNEX,  $F < 1$  (n.s).

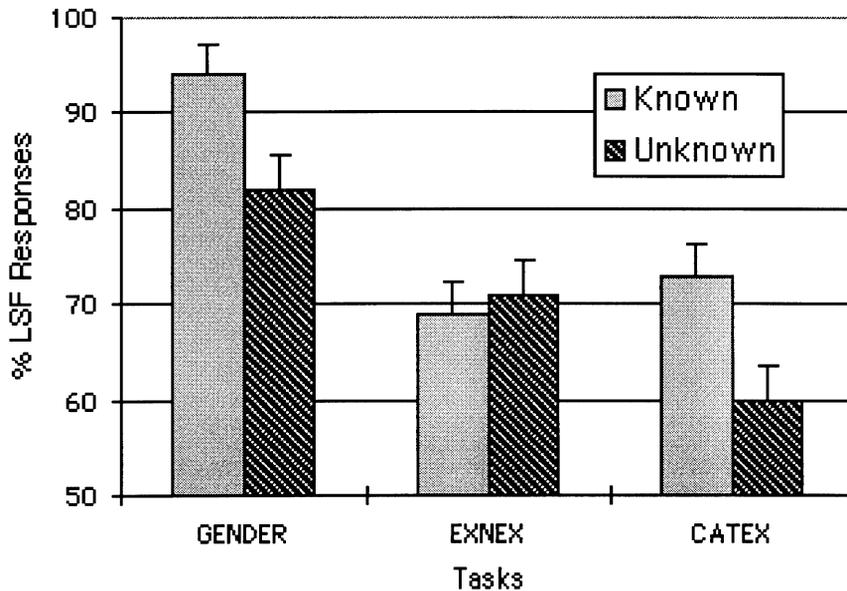


Fig. 4. The scale biases observed in Experiment 3. The data illustrate that all groups were biased to LSF. However, those who knew the faces that comprised the test hybrids were more biased to LSF in the GENDER and CATEX conditions, but not in EXNEX.

It could be argued that the strengthening of the LSF bias in K\_GENDER and K\_CATEX does not arise from a knowledge of face identities *per se*, but from using the same faces in training and testing (a form of repetition priming). If repetition was the only factor at play, we should expect K\_EXNEX to be similarly enhanced, but it was not. The strengthening in K\_GENDER is easy to understand because identity discloses gender. The absence of a strengthening in K\_EXNEX remains to be explained. For CATEX, the observed influence of identity on expression does not imply a serial processing of the former before the latter. Instead, the influence only implies that judgments of identity and expressions are not *independent*. In a Garner interference paradigm (Garner, 1974), Schweinberger and Soukup (in press) recently found that identity influenced judgments of expression whereas expression did not influence identity judgments. Shortly put, these two processes were in a relation of asymmetric influence, revealing that they were interactive, not computationally encapsulated. Likewise, our results showed an influence of identity on the scale bias of judgments of expressions.

In sum, the first identification task induced a LSF perceptual bias which then transferred across tasks to EXNEX, CATEX and GENDER, with a marked strengthening when faces were known. As in Experiment 2, the transfer of bias across tasks revealed a lasting perceptual effect, confirming the cognitive determination of perception reported earlier.

## 5. General discussion

From an empirical standpoint, we wanted to demonstrate the flexible use of scale information for different face categorization tasks. In Experiment 1, an expressive vs. non-expressive (EXNEX) task was biased to HSF, a categorization of the expression itself (CATEX) was biased to LSF, but both coarse and fine scales allowed gender decisions (GENDER). However, these biases were not fixed. Examining the dynamics of categorization-induced scale biases in Experiment 2 revealed that the bias of a first task (LSF from CATEX vs. HSF from EXNEX) transferred to the GENDER so that different groups tended to see the opposite gender in the same stimulus. The initial learning of face identity in Experiment 3 was found to be biased to LSF. Testing for the transfer of bias in GENDER, EXNEX and CATEX with known vs. unknown faces revealed that all tasks became LSF-biased. This bias was stronger for GENDER and CATEX (but not for EXNEX) when the faces were known. Thus, the evidence suggests that the diagnosticity of scale cues in a task, together with perceptual set effects, best predict scale use in the categorizations tested here.

We should be careful and emphasize that the reported biases might not all apply to the normal perception of faces. Our face stimuli were ambiguous, normalized, presented tachistoscopically and designed to capture flexible use and perception of spatial scales. However, hybrids still represented faces, and they were perceived as faces whether subjects used the high or low frequency information. This flexible use raises a number of issues that should be of interest to face recognition research, including the specificity of scale cues for various face detection and categorization tasks; the development of (global or local, Oliva and Schyns, 1997) scale cues with the acquisition of face expertise (see Tanaka and Gauthier, 1997 for discussions); and the dichotomy between sufficient scale information vs. its use in natural recognition tasks. Although these research topics apply directly to faces, it is worth stressing that flexible perceptions of hybrid scenes were obtained in Oliva and Schyns (1997), albeit via other means. It is therefore an interesting empirical issue to examine whether differences in categorization tasks (e.g. a basic-level, *city*, vs. a subordinate, *New York* categorization of the same picture of New York) can also determine flexible scale use in scenes and other objects. From our results, we would predict that this depends on whether these categorizations require diagnostic, scale-specific cues (Schyns, 1998).

One could object that casual observation of hybrids reveals only their HSF cues, not their LSF cues (see Fig. 1). One could therefore ask whether LSF information really contributes to naturalistic recognition, or whether its use is limited to the kind of tachistoscopic studies presented here. Three points should be made. First, given that all spatial frequencies are integrated early on, cues at different scales *are* available to recognize the visual input. Secondly, recognition is not limited to foveal vision (where receptor density allows the representation of HSF luminance changes) and we know from eye movement research that a good deal of naturalistic recognition takes place in the periphery, using LSF information. Everyday recognition also needs LSF to detect faces (Morrisson and Schyns, in press), to recognize them in

naturalistic conditions such as a smoky environment, low contrast, through a dirty window, and so forth. Thirdly, research in image compression has repeatedly shown that HSF might not always be as important as one would intuitively think. An image can be computed that looks identical to the original (i.e. it takes several seconds to notice a difference) even though it misses significant HSF information in selective places (see Strang and Nguyen, 1997 for examples). Image compression algorithms (which are ubiquitous in modern computing) are based on the fact that LSF is the skeleton of an image that HSF only fleshes out. For some portions of the image, the skeleton suffices. In sum, there could be more to LSF than meets the eye.

### 5.1. *Methodological implications*

From a methodological standpoint, the results reported here emphasize the importance of studying many (not just one) categorizations of an identical face, object, or scene before inferring its memory representation. The study of a single categorization task might not be sufficient to tap into face, object or scene representations *per se*. To illustrate, if it was discovered that the information demands of one categorization task were X (e.g. LSF or HSF cues), then it would be straightforward to assume that the representation of the face, object or scene was effectively that X. However, how would we know whether X represents the object, or the task itself? This point is not just methodological, and it might be the most important implication of our results for recognition studies (see also Schyns, 1998).

On a pessimistic note, our results could lead to the conclusion that the factors of the exact nature of the stimuli, of the task, of set effects, and so forth interact in a complex way so that one cannot conclude that there is a certain bias/representation for certain types of category judgments. For example, we observed opposite biases for the same GENDER task when subjects first solved CATEX or EXNEX. We also showed a reversal of bias from HSF to LSF when subjects who had learned the identity of faces performed EXNEX. We pointed out that the LSF bias observed for the identification of a few faces might turn into an HSF bias if the task included many more faces. Thus, one could conclude from the actual results that scale selection is just too flexible to observe consistent, categorization-specific biases.

As already pointed out, our methods sought to find evidence of flexibility. More naturalistic conditions might elicit more consistency. For example, whereas our stimuli normalized the important hairstyle, its distribution in natural images could trigger a systematic LSF bias when assessing the gender of faces. However, to the extent that different spatial scales do support different sorts of object cues, changes of biases should occur together with the development of object representations (see Christensen et al., 1981; Biederman and Shiffrar, 1987; Tanaka and Taylor, 1991; Norman et al., 1992; Schyns and Murphy, 1994; Gauthier and Tarr, 1997; Schyns and Rodet, 1997; Shiffrin and Lightfoot, 1997;). Entertaining the possibility of representational and perceptual changes throughout development does not facilitate object recognition research (because different representations of the same object could accompany different levels of expertise). However, it opens interesting

research questions about the processes whereby perceptual systems optimize object representations to their own circumstances.

### 5.2. *Theoretical implications*

From a theoretical standpoint, an interesting observation of studies with hybrids is that their immediate perceptual appearance changes with the categorization task. The evidence so far is indirect. In all three experiments, debriefing revealed that many subjects who categorized at their diagnostic scale were unaware of relevant information present at the other scale. Furthermore, most subjects expressed surprise when they were informed of the stimulus composition. Future research should unravel the precise influence the categorization task exerts on the perception of a face, object, or scene. One promising avenue arises from the common underpinnings of hybrid stimuli and the spatial filtering techniques that are ubiquitous in the psychophysics of early vision. This enables a study of hybrid recognition in conjunction with psychophysical techniques to understand whether attention to a diagnostic spatial scale (or neglect of the other scale) affects the filtering properties (e.g. contrast thresholds, orientation selectivity) of the earliest stages of visual processing. Evidence that it does would have far-reaching implications for classical issues in cognitive science ranging from the depth of feedback loops in early vision, the early vs. late selection models of attention (He et al., 1996), the bi-directionality of cognition (Schyns, 1997), the sparse vs. exhaustive perceptions of distal stimuli (Hochberg, 1982), to the cognitive penetrability of vision (Fodor, 1983).

For example, the idea of the continuity between perception and cognition defended here (see also Bruner, 1957; Schyns et al., 1998) is challenged in Pylyshyn's (in press) recent assertion that early vision is by and large cognitively impenetrable. However, Pylyshyn notes that one exception arises when attention must be allocated to certain locations or certain object properties prior to the operation of early vision. The diagnostic allocation of attention to the content of spatial scales reported here therefore satisfies to the condition for cognitive penetrability of Pylyshyn (in press). However, it would seem that the 'allocation of attention to visual properties prior to the operation of early vision' does not reduce cognitive penetrability to a few situations of recognition. Instead, the most common situations concern the basic (*bird, car*) vs. subordinate (*sparrow, Mercedes*) categorizations of objects which are known to require different cues from the visual input (as LSF and HSF cues were best-suited to different face categorizations in the experiments just described). It is therefore conceivable that the common basic vs. subordinate categorizations of identical objects would elicit distinct perceptual settings of early vision. Cognitive studies with hybrids, together with the psychophysical testing of early vision, could offer a powerful platform to investigate further this complicated issue.

## 6. **Concluding remarks**

The studies reported here examined the possibility that different categorizations

of the same faces could change their perception. Specifically, we tested whether the categorizations elicited a differential use of the spatial scales supporting the perception of the input. The evidence that they do raise a number of new issues as to how closely vision is bound to cognition.

### Acknowledgements

The authors wish to thank Paula Niedenthal from the Psychology Department at Indiana University for lending us the original face stimuli that were used in our experiments for computing hybrids. We also wish to thank Pierre Demartines and Jeanny Hérault from LTIRF at INPG, Grenoble, France, for useful technical help with stimulus computation. Many thanks to Frederic Gosselin, Gregory L. Murphy, Pascal Mammassian, Donald Morrisson and three anonymous reviewers for helpful comments which have greatly improved the quality of this manuscript. This work was partially funded by ESRC grant R000237412 awarded to Philippe G. Schyns and a Fyssen Foundation postdoctoral research fellowship (Paris, France) awarded to Aude Oliva.

### References

- Bachmann, T., 1991. Identification of spatially quantised tachistoscopic images of faces: how many pixels does it take to carry identity? *European Journal of Cognitive Psychology* 3, 85–103.
- Bachmann, T., 1997. The effect of coarseness of quantisation, exposure duration, and selective spatial attention on the perception of spatial quantized ('blocked') visual images. *Perception* 26, 1181–1196.
- Barrera, M., Maurer, D., 1983. The perception of facial expression by 3-month-olds. *Child Development* 52, 203–206.
- Bassili, J.N., 1979. Emotion recognition: the role of facial movement and the relative importance of upper and lower areas of the face. *Journal of Personality and Social Psychology* 37, 2049–2059.
- Biederman, I., Shiffrar, M.M., 1987. Sexing day-old chicks: A case study and expert systems analysis of a difficult perceptual-learning task. *Journal of Experimental Psychology: Learning, Memory and Cognition* 13, 640–645.
- Blackemore, C., Campbell, F.W., 1969. On the existence of neurons in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology (London)* 203, 237–260.
- Breitmeyer, B.G., Ganz, L., 1976. Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression and information processing. *Psychological Review* 83, 1–35.
- Bruce, V., 1994. What the human face tells the human mind: Some challenges for the robot-human interface. *Advanced Robotics* 8, 341–355.
- Bruner, J.S., 1957. On perceptual readiness. *Psychological Review* 64, 123–152.
- Calder, A.J., Young, A.W., Perrett, D.I., Etcoff, N.L., Rowland, D., 1996. Categorical perception of morphed facial expressions. *Visual Cognition* 3, 81–117.
- Campbell, F.W., Robson, J.G., 1968. Application of the Fourier analysis to the visibility of gratings. *Journal of Physiology (London)* 88, 551–556.
- Carey, S., 1992. Becoming a face expert. *Philosophical Transactions of the Royal Society of London B355*, 95–103.

- Christensen, E.E., Murry, R.C., Holland, K., Reynolds, J., Landay, M.J., Moore, J.G., 1981. The effect of search time on perception. *Radiology* 138, 361–365.
- Costen, N.P., Parker, D.M., Craw, I., 1994. Spatial content and spatial quantisation effects in face recognition. *Perception* 23, 129–146.
- Costen, N.P., Parker, D.M., Craw, I., 1996. Effects of high-pass and low-pass spatial filtering on face identification. *Perception and Psychophysics* 38, 602–612.
- Dailey, M.N., Cottrell, G.W., (in press). Task and spatial frequency effects on face specialization. *Advances in Neural Information Processing Systems*, Vol. 10. MIT Press, Cambridge, MA.
- De Valois, R.L., De Valois, K.K., 1990. *Spatial Vision*. Oxford University Press, New York.
- Diamond, R., Carey, S., 1990. On the acquisition of pattern encoding skills. *Cognitive Development* 5, 345–368.
- Eckman, P., Friesen, W.V., 1975. *Pictures of Facial Effects*. Consulting Psychologists Press, Palo Alto, CA.
- Etcoff, N.L., Maggee, J.J., 1992. Categorical perception of facial expressions. *Cognition* 44, 227–240.
- Findlay, J.M., Brogan, D., Wenban-Smith, M., 1993. The visual signal for saccadic eye movements emphasizes visual boundaries. *Perception and Psychophysics* 53, 633–641.
- Fodor, J., 1983. *The Modularity of Mind*. MIT Press.
- Garner, 1974. *The Processing of Information and Structure*. Erlbaum, Potomac, MD.
- Gauthier, I., Tarr, M.J., 1997. Becoming a ‘greeble’ expert: exploring mechanisms for face recognition. *Vision Research* 37, 1673–1682.
- Gray, M.S., Lawrence, D.T., Golomb, B.A., Sejnowski, T.J., 1995. A perceptron reveals the face of sex. *Neural Computation* 7, 1160–1164.
- Hammett, S.T., Snowden, R.J., 1995. The effect of contrast adaptation on briefly presented stimuli. *Vision Research* 35, 1721–1725.
- Harmon, L.D., 1973. The recognition of faces. *Scientific American* 229, 71–82.
- He, S., Cavanagh, P., Intriligator, J., 1996. Attentional resolution and the locus of visual awareness. *Nature* 383, 334–337.
- Henning, G.B., Hertz, B.G., Broadbent, D.E., 1975. Some experiments bearing on the hypothesis that the visual system analyzes spatial patterns in independent bands of spatial frequency. *Vision Research* 15, 887–899.
- Hochberg, J., 1982. How big is a stimulus? In: Beck, J. (Ed.), *Organization and Representation in Perception*. Lawrence Erlbaum, Hillsdale, NJ, pp. 191–217.
- Hughes, H.C., Nozawa, G., Kitterle, F., 1996. Global precedence, spatial frequency channels, and the statistics of natural images. *Journal of Cognitive Neuroscience* 8, 197–230.
- Jacobs, R., Jordan, M., Nowlan, S., Hinton, G., 1991. Adaptive mixtures of local experts. *Neural Computation* 3, 79–87.
- Jenkins, J., Craven, B., Bruce, V., Akamatsu, S., 1997. Methods for detecting social signals from the face. Technical Report of IECE, HIP96-39, The Institute of Electronics, Information and Communication Engineers, Japan.
- Kuchuk, A., Vibbert, M., Bornstein, M.H., 1986. The perception of smiling and its experiential correlates in 3-month-old infants. *Child Development* 57, 1054–1061.
- Legge, G.E., Gu, Y., 1989. Stereopsis and contrast. *Vision Research* 29, 989–1004.
- Long, G., Toppino, T.C., Mondin, G.W., 1992. Prime time: fatigue and set effects in the perception of reversible figures. *Perception and Psychophysics* 52, 609–616.
- Marr, D., Hildreth, E.C., 1980. Theory of edge detection. *Proceedings of the Royal Society of London Series B* 207, 187–217.
- Marshall, J.A., Burbeck, C.A., Ariely, J.P., Rolland, J.P., Martin, K.E., 1996. *Journal of the Optical Society of America A* 13, 681–688.
- Mayer, R.E., 1992. *Thinking, Problem Solving, Cognition*. Freeman, New York.
- Morgan, M.J., 1992. Spatial filtering precedes motion detection. *Nature* 355, 344–346.
- Morrison, D., Schyns, P.G., (in press). Exploring the interaction between face processing and attention. *Perception*.

- Norman, G.R., Brooks, L.R., Coblenz, C.L., Babcock, C.J., 1992. The correlation of feature identification and category judgments in diagnostic radiology. *Memory and Cognition* 20, 344–355.
- Oliva, A., Schyns, P.G., 1995. Mandatory scale perception promotes flexible scene categorization. *Proceedings of the XVII Meeting of the Cognitive Science Society*. Lawrence Erlbaum, Hillsdale, NJ, pp. 159–163.
- Oliva, A., Schyns, P.G., 1997. Coarse blobs or fine edges? Evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive Psychology* 34, 72–107.
- Pantle, A., Sekuler, R., 1968. Size detecting mechanisms in human vision. *Science* 162, 1146–1148.
- Parker, D.M., Lishman, J.R., Hughes, J., 1992. Temporal integration of spatially filtered visual images. *Perception* 21, 147–160.
- Parker, D.M., Lishman, J.R., Hughes, J., 1996. Role of coarse and fine information in fact and object processing. *Journal of Experimental Psychology: Human Perception and Performance* 22, 1448–1466.
- Pylyshyn, Z., (in press). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*.
- Rosenblum, M., Yacoob, Y., Davis, L.S., 1996. Human expression recognition from motion using a radial function network architecture. *IEEE Transactions on Neural Networks* 7, 1121–1138.
- Schweinberger, S.R., Soukup, G.R., (in press). Asymmetric relationships among the perception of facial identity, emotion and facial speech. *Journal of Experimental Psychology: Human Perception and Performance*.
- Schyns, P.G., 1997. Categories and percepts: a bi-directional framework for categorization. *Trends in Cognitive Sciences* 1, 183–189.
- Schyns, P.G., 1998. Diagnostic recognition: task constraints, object information and their interactions. *Cognition* 67, 147–179.
- Schyns, P.G., Goldstone, R.L., Thibaut, J.P., 1998. The development of features in object concepts. *Behavioral and Brain Sciences* 21, 17–41.
- Schyns, P.G., Murphy, G.L., 1994. The ontogeny of part representation in object concepts. In: Medin, D.E. (Ed.), *The Psychology of Learning and Motivation*. Academic Press, San Diego, CA, Vol. 31, pp. 301–349.
- Schyns, P.G., Oliva, A., 1994. From blobs to boundary edges: Evidence for time and spatial scale dependent scene recognition. *Psychological Science* 5, 195–200.
- Schyns, P.G., Rodet, L., 1997. Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory and Cognition* 23, 681–696.
- Sergent, J., 1986. Microgenesis of face perception. In: Ellis, D.H., Jeeves, M.A., Newcombe, F., Young, A. (Eds.), *Aspects of Face Processing*. Martinus Nijhoff, Dordrecht, pp. 17–73.
- Sergent, J., Ohta, S., MacDonald, B., 1992. Functional neuroanatomy of face and object processing: a positron emission tomography study. *Brain* 115, 15–36.
- Shiffrin, R.M. and Lightfoot, N., 1997. Perceptual learning of alphanumeric-like characters. In: Goldstone, R., Schyns, P.G., Medin, D.E. (Eds.), *Mechanisms of Perceptual Learning*. Academic Press, San Diego, CA, pp. 45–80.
- Snowden, R.J., Hammett, S.T., 1992. Subtractive and divisive adaptation in the human visual system. *Nature* 355, 248–250.
- Strang, G., Nguyen, T., 1997. *Wavelets and Filter Banks*. Wellesley-Cambridge Press, Wellesley, MA.
- Tanaka, J., Gauthier, I., 1997. Expertise in object and face recognition. In: Goldstone, R., Schyns, P.G., Medin, D.E. (Eds.), *Mechanisms of Perceptual Learning*. Academic Press, San Diego, CA, pp. 85–121.
- Tanaka, J., Farah, M.J., 1993. Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology* 46A, 225–245.
- Tanaka, J., Sengco, J.A., 1997. Features and their configuration in face recognition. *Memory and Cognition* 25, 583–592.
- Tanaka, J., Taylor, M.E., 1991. Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology* 15, 121–149.